CLAIMS

1. A processing device programmed to determine homing paths for a plurality of virtual private local area network services in an Ethernet network comprising a plurality of PE nodes, programmed to perform the steps of:

computing a plurality of sets of different homing configurations;

wherein each homing configuration in each set of different homing configurations is computed by a respective iteration of steps;

wherein each iteration corresponds to a respective virtual private local area network service in the plurality of virtual private local area network services and for a respective selected layer two provider edge node in the Ethernet network; and

wherein each iteration comprises the steps of:

selecting an ingress PE node and an egress PE node;

determining bandwidth into the ingress PE node;

determining bandwidth out of the egress PE node;

specifying a first path for communication from the ingress PE node to the egress PE node and a second path for communication from the egress PE node to the ingress PE node, wherein each path of the first and second paths comprises at least one P node;

computing a cost function for each set of different homing configurations in the plurality of sets of different homing configurations; and

selecting a set of homing configurations from the plurality of sets of different homing configurations in response to a respective computed cost function.

2. The processing device of claim 1:

wherein each homing configuration comprises primary homing paths and secondary homing paths;

wherein the primary homing paths comprise the first and second paths;

and wherein the processing device is further programmed to determine the secondary homing paths for each homing configuration, wherein each iteration further comprises the steps of:

selecting a secondary ingress PE node and a secondary egress PE node;

determining bandwidth into the secondary ingress PE node; and

determining bandwidth out of the secondary egress PE node;

specifying a third path for communication from the secondary ingress PE node to the secondary egress PE node and a fourth path for communication from the secondary egress PE node to the secondary ingress PE node, wherein each path of the third and fourth paths comprises at least one P node.

3. The processing device of claim 2 wherein each iteration further comprises the steps of:

determining whether either the first path or the second path violates a constraint set comprising at least one constraint; and

determining whether either the third path or the forth path violates the constraint set.

4. The processing device of claim 2 wherein the constraint set comprises a constraint relating to the bandwidth along the first and second paths.

5. The processing device of claim 2 wherein the constraint set comprises a constraint selected from a set consisting of:

a first constraint wherein a layer two provider edge node may multi-home only to PE nodes within a predetermined geographical proximity of the layer two provide edge node;

a second constraint wherein a logical VPLS entity may connect to only one PE node;

a third constraint, wherein traffic between a logical VPLS entity and a PE node is less than or equal to available bandwidth on a link between the logical VPLS entity and the PE node;

a fourth constraint, wherein bandwidth is equalized in that a sum of bandwidth entering a PE node plus a sum of bandwidth entering that PE node from any P nodes equals a sum of bandwidth leaving that PE node plus a sum of bandwidth leaving that PE node towards any P nodes;

a fifth constraint, wherein at a P node, a sum of bandwidth entering that P node must be less than or equal to the sum of the bandwidth capabilities on the output link between that P node and an adjacent P node;

a sixth constraint, wherein a sum of bandwidth entering all PE nodes must equal a sum of bandwidth leaving all PE nodes;

a seventh constraint, wherein bandwidth into a PE node may not exceed any limit provided by the processing capability of that PE node; and

an eighth constraint, wherein end to end delay limits are provided for one or more paths.

6. The processing device of claim 2 wherein the plurality of PE nodes comprise fully-meshed PE nodes.

7. The processing device of claim 2 and further programmed to perform the steps of, prior to the computing step:

inputting a set of P nodes in the Ethernet network; and

inputting bandwidth capability between each P node, in the set of P nodes, and an adjacent P node in the set of P nodes.

8. The processing device of claim 7 and further programmed to perform the steps of, prior to the computing step:

inputting a set of layer two provider edge nodes in the Ethernet network; and

inputting bandwidth capability between each layer two provide edge node, in the set of layer two provider edge nodes, and each PE node in the plurality of PE nodes.

9. The processing device of claim 8 and further programmed to perform the step of, prior to the computing step, inputting bandwidth requirements information as between each customer edge node connected to the Ethernet network and any other customer edge node connected to the Ethernet network.

10. The processing device of claim 2 wherein the step of specifying a third path and a fourth path is in response to a number of P node hops in the third path and the fourth path.

11. The processing device of claim 2 wherein the step of specifying a third path and a fourth path is in response to distributing bandwidth load along a number of P node hops in the third path and the fourth path such that bandwidth load between each successive hop between successive P nodes is within a tolerance of a bandwidth load between each other successive hop between successive P nodes in a same path.

12. The processing device of claim 2 wherein the step of specifying a first path and a second path is in response to a number of P node hops in the first path and the second path.

13. The processing device of claim 2 wherein the step of specifying a first path and a second path is in response to distributing bandwidth load along a number of P node hops in the first path and the second path such that bandwidth load between each successive hop between successive P nodes is within a tolerance of a bandwidth load between each other successive hop between successive P nodes in a same path.

14. The processing device of claim 2 wherein the step of computing a cost function for each set of different homing configurations in the plurality of sets of different homing configurations comprises computing a total number of virtual private local area network service connections included in each set of different homing configurations.

15. The processing device of claim 14 wherein the step of selecting a set of homing configurations comprises selecting a set of homing configurations having a maximum total number of virtual private local area network service connections included in the selected set of homing configurations.

16. The processing device of claim 1 wherein the step of computing a cost function for each set of different homing configurations in the plurality of sets of different homing configurations comprises computing a total number of virtual private local area network service connections included in each set of different homing configurations.

17. The processing device of claim 16 wherein the step of selecting a set of homing configurations comprises selecting a set of homing configurations having a maximum total number of virtual private local area network service connections included in the selected set of homing configurations.

18. The processing device of claim 1 wherein the network comprises a Metro Ethernet network.

19. The processing device of claim 1 wherein one node in the plurality of nodes comprises the processing device.

20. The processing device of claim 1 and further programmed to perform the step of communicating information to PE nodes in the Ethernet network to configure nodes receiving the information, wherein the information is in response to the selected set of homing configurations.

21. A method of electronically determining homing paths for a plurality of virtual private local area network services in an Ethernet network comprising a plurality of PE nodes, comprising the steps of:

determining a plurality of sets of different homing configurations;

wherein each homing configuration in each set of different homing configurations is computed by a respective iteration of steps;

wherein each iteration corresponds to a respective virtual private local area network service in the plurality of virtual private local area network services and for a respective selected layer two provider edge node in the Ethernet network; and

wherein each iteration comprises the steps of:

selecting an ingress PE node and an egress PE node;

determining bandwidth into the ingress PE node;

determining bandwidth out of the egress PE node;

specifying a first path for communication from the ingress PE node to the egress PE node and a second path for communication from the egress PE node to the ingress PE node, wherein each path of the first and second paths comprises at least one P node;

determining a cost function for each set of different homing configurations in the plurality of sets of different homing configurations; and

selecting a set of homing configurations from the plurality of sets of different homing configurations in response to a respective computed cost function.

* * * * *